arXiv:1612.03350v1 [cs.CE] 10 Dec 2016

Non-negative Factorization of the Occurrence Tensor from Financial Contracts

Zheng Xu[1], Furong Huang[3], Louiqa Raschid[2], Tom Goldstein[1]
[1]Department of Computer Science, [2]Smith School of Business, University of Maryland
[3]Microsoft Research, New York City

Abstract

We propose an algorithm for the non-negative factorization of an occurrence tensor built from heterogeneous networks. We use ℓ_0 norm to model sparse errors over discrete values (occurrences), and use decomposed factors to model the embedded groups of nodes. An efficient splitting method is developed to optimize the nonconvex and nonsmooth objective. We study both synthetic problems and a new dataset built from financial documents, resMBS.

1 Introduction

Tensor factorization is a powerful approach to a myriad of unsupervised learning problems [1, 14, 17]. We propose an efficient algorithm called non-negative occurrence tensor factorization (NOTF) for analyzing heterogeneous networks. As an application of NOTF, we study the resMBS dataset[5, 19], which consists of the relationships that financial institutions (FIs, e.g., Bank of America) play roles (e.g., issuer) in financial contracts (FCs). ResMBS is automatically extracted from a collection of public financial contracts . We can represent resMBS as a three-mode (FC, FI, Role) occurrence tensor with non-negative discrete values corresponding to the confidence of the relationship extraction.

The occurrence tensor has several characteristics. First, the tensor has positive discrete values. Second, the observed tensor contains sparse noise corresponding to errors caused by inaccurate extraction. Most importantly, we expect that FIs will play specific roles across multiple FCs to create a community. The observed occurrence tensor could then be decomposed into a low-rank tensor together with sparse noise. This decomposition can be used to clean and complete the observation, and also allow us to understand the underlying (FC, FI, Role) communities.

The proposed NOTF extends the CANDECOMP/PARAFAC (CP) decomposition [6, 12] of a tensor. After decomposing the occurrence tensor, each rank-one tensor component represents a (FC, FI, Role) community. Hence the decomposed factors are constrained to be non-negative. Instead of the ℓ_2 norm for factorization of real valued tensors with Gaussian noise, the ℓ_0 norm is considered for the discrete values and the sparse errors. We develop an algorithm based on the alternating direction method of multipliers (ADMM) [4, 22] to optimize the nonconvex and nonsmooth objective of NOTF.

We briefly discuss some related research. While tensor decomposition dates back to the 1920s, robust tensor decomposition solutions have been presented in some recent papers [14, 2, 8, 11]. Standard CP decomposition has been applied to heterogenous networks [16]. Discrete valued tensors have been studied in [18]. ADMM has been extensively used as a solver for robust tensor recovery [8, 11], and also for standard non-negative tensor factorization[15]. Non-ℓ_2 norms have been used for CP decomposition of real valued tensors in computer vision [13, 7]. ADMM for ℓ_0 norm was empirically studied in [21]. Probabilistic communities were discussed for financial documents in [19].

2 Non-negative Occurrence Tensor Factorization

We use notations similar to [14]. Vectors and matrices are denoted by lowercase and capital letters, respectively. Higher mode tensors are denoted by Euler script letters. Three mode tensors $\mathcal{X} \in \mathbb{R}^{N_1 \times N_2 \times N_3}$ are used as an example in this paper. Fibers are column vectors extracted from tensors by fixing every index but one, e.g., $x_{:jk}, x_{i:k}, x_{ij:}$. Mode-d fibers are arranged to get the mode-d

unfolding matrix $X_{(d)}$, e.g., $X_{(1)} = [x_{:jk}] \in \mathbb{R}^{N_1 \times N_2 N_3}$. Denote the vector outer product by $\circ$, then the CP decomposition factorizes a tensor into a sum of rank-one tensors as $\mathcal{X} = \sum_{r=1}^{R} a_r \circ b_r \circ c_r$. If we denote $A = [a_r] \in \mathbb{R}^{N_1 \times R}$, $B = [b_r] \in \mathbb{R}^{N_2 \times R}$ and $C = [c_r] \in \mathbb{R}^{N_3 \times R}$, the Kronecker product $A \otimes B \in \mathbb{R}^{N_1 N_2 \times RR}$, and the Khatri-Rao product $A \odot B = [a_r \otimes b_r] \in \mathbb{R}^{N_1 N_2 \times R}$, then we can compactly represent CP as $\mathcal{X} = \sum_{r=1}^{R} a_r \circ b_r \circ c_r = [\![A, B, C]\!]$, which is equivalent to the unfolding representation $X_{(1)} = A(C \odot B)^T$, $X_{(2)} = B(C \odot A)^T$, and $X_{(3)} = C(B \odot A)^T$.

We are seeking to recover a tensor $\mathcal{X} = [\![A, B, C]\!]$ that is at most rank R from the noisy observations $\mathcal{O}$. Each rank-one tensor $a_r \circ b_r \circ c_r$ captures a community and $a_r \geq 0, b_r \geq 0, c_r \geq 0$ represent the weights of nodes (e.g., FI, FC, and roles) in the community. For discrete valued tensors, we minimize the sparse error measured with the ℓ_0 norm rather than the ℓ_2 loss in standard decomposition,

$$\min_{A,B,C} \|[\![A, B, C]\!] - \mathcal{O}\|_0, \text{ subject to } A \geq 0, B \geq 0, C \geq 0, \tag{1}$$

where $\|\mathcal{X}\|_0 = \sum_{i,j,k} \mathbb{1}_{\{z:\,|z|>0\}}(x_{ijk})$ is the counts of nonzero values in a tensor, $\mathbb{1}_S$ is the indicator function of the set S: $\mathbb{1}_S(v) = 1$, if $v \in S$, and $\mathbb{1}_S(v) = 0$, otherwise.

We minimize the NOTF objective (1) by introducing an intermediate variable $\mathcal{U}$,

$$\min_{\mathcal{U},A,B,C} \|\mathcal{U}\|_0 + \iota_{\{z:\,z\geq 0\}}(A, B, C), \text{ subject to } \mathcal{U} = [\![A, B, C]\!] - \mathcal{O}, \tag{2}$$

where ι_S is the characteristic function of the set S; $\iota_S(v) = 0$, if $v \in S$, and $\iota_S(v) = \infty$, otherwise. We then apply alternating direction method of multipliers (ADMM) [9, 4, 10, 22] by introducing dual variables λ and alternatively solving subproblems of $\mathcal{U}$ and A, B, C, with p indexing iterations,

$$\mathcal{U}^{p+1} = \arg\min_{\mathcal{U}} \|\mathcal{U}\|_0 + \frac{\tau}{2}\left\|\mathcal{U} - [\![A^p, B^p, C^p]\!] + \mathcal{O} + \lambda^p\right\|_F^2 \tag{3}$$

$$A^{p+1}, B^{p+1}, C^{p+1} = \arg\min_{A,B,C} \iota_{\{z:\,z\geq 0\}}(A, B, C) + \frac{\tau}{2}\left\|\mathcal{U}^{p+1} - [\![A, B, C]\!] + \mathcal{O} + \lambda^p\right\|_F^2 \tag{4}$$

$$\lambda^{p+1} = \lambda^p + \mathcal{U}^{p+1} - [\![A^{p+1}, B^{p+1}, C^{p+1}]\!] + \mathcal{O}, \tag{5}$$

where $\|\mathcal{X}\|_F = \sqrt{\sum_{i,j,k} x_{ijk}^2}$, and τ is a hyperparameter called the penalty parameter.

Subproblem (3) can be solved by the proximal operator of the ℓ_0 norm, known as hard-thresholding,

$$\mathcal{U}^{p+1} = \text{hard}([\![A^p, B^p, C^p]\!] - \mathcal{O} - \lambda^p,\, 1/\tau), \tag{6}$$

where $\text{hard}(\mathcal{Z}, t) = \arg\min_{\mathcal{X}} \|\mathcal{X}\|_0 + 1/2t\|\mathcal{X} - \mathcal{Z}\|_F^2 = \mathcal{Z} * \mathcal{I}_{\{z:|z|>\sqrt{2t}\}}(\mathcal{Z})$, with $*$ representing the element-wise Hadamard product, and $\mathcal{I}_S(\mathcal{X}) = [\mathbb{1}_S(x_{ijk})]$ r the element-wise indicator function.

Subproblem (4) is nonnegative tensor factorization, which can be solved by alternatively optimizing one of A, B, or C when the other two are fixed, with q indexing iterations,

$$A^{p,q+1} = \arg\min_{A} \iota_{\{z:\,z\geq 0\}}(A) + \frac{\tau}{2}\|(\mathcal{U}^{p+1} + \mathcal{O} + \lambda^p)_{(1)} - A(C^{p,q} \odot B^{p,q})^T\|_F^2 \tag{7}$$

$$= \max\{(\mathcal{U}^{p+1} + \mathcal{O} + \lambda^p)_{(1)}(C^{p,q} \odot B^{p,q})((C^{p,q})^T(C^{p,q}) * (B^{p,q})^T(B^{p,q}))^\dagger,\, 0\} \tag{8}$$

$$B^{p,q+1} = \arg\min_{B} \iota_{\{z:\,z\geq 0\}}(B) + \frac{\tau}{2}\|(\mathcal{U}^{p+1} + \mathcal{O} + \lambda^p)_{(2)} - B(C^{p,q} \odot A^{p,q+1})^T\|_F^2 \tag{9}$$

$$= \max\{(\mathcal{U}^{p+1} + \mathcal{O} + \lambda^p)_{(2)}(C^{p,q} \odot A^{p,q+1})((C^{p,q})^T(C^{p,q}) * (A^{p,q+1})^T(A^{p,q+1}))^\dagger,\, 0\} \tag{10}$$

$$C^{p,q+1} = \arg\min_{C} \iota_{\{z:\,z\geq 0\}}(C) + \frac{\tau}{2}\|(\mathcal{U}^{p+1} + \mathcal{O} + \lambda^p)_{(3)} - C(B^{p,q+1} \odot A^{p,q+1})^T\|_F^2 \tag{11}$$

$$= \max\{(\mathcal{U}^{p+1} + \mathcal{O} + \lambda^p)_{(3)}(B^{p,q+1} \odot A^{p,q+1})((B^{p,q+1})^T(B^{p,q+1}) * (A^{p,q+1})^T(A^{p,q+1}))^\dagger,\, 0\}. \tag{12}$$

Each subproblem is a constrained least squares problem the recovers the mode-d unfolding matrix $(\mathcal{U}^{p+1} + \mathcal{O} + \lambda^p)$, starting from $A^{p,0} = A^p, B^{p,0} = B^p, C^{p,0} = C^p$, and updating A, B, C according to (7)-(12) until convergence, then $A^{p+1} = A^{p,\text{end}}, B^{p+1} = B^{p,\text{end}}, C^{p+1} = C^{p,\text{end}}$. The updates (7)-(12) for the subproblems (4) usually converge in less than ten iterations when warm started from the previous iteration. Relative "residuals" are used to monitor the convergence of (3)-(5), and are defined by

$$\text{res}_1 = \frac{\left\|[\![A^{p,q}, B^{p,q}, C^{p,q}]\!] - [\![A^{p,q-1}, B^{p,q-1}, C^{p,q-1}]\!]\right\|_F}{\left\|[\![A^{p,q-1}, B^{p,q-1}, C^{p,q-1}]\!]\right\|_F} \tag{13}$$

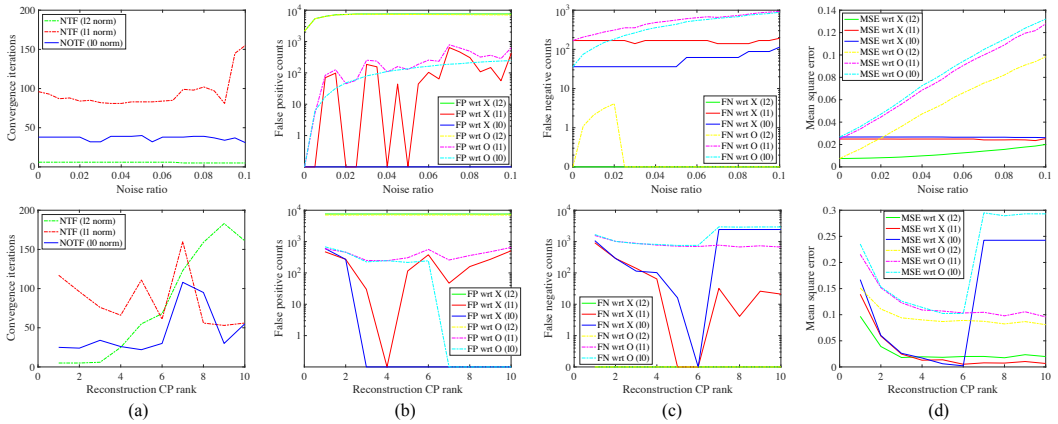

Figure 1: (a) Convergence iteration, (b) false positive count, (c) false negative count, and (d) mean square error when varying the noise ratio (top) and varying the tensor CP rank R for NOTF (bottom) for the synthetic dataset. We set rank $R = 3$ when varying the noise ratio (top) and set a noise ratio of 10% when varying the CP rank (bottom). Note that reconstruction errors with respect to both groundtruth $\mathcal{X}$ and noisy observation $\mathcal{O}$ are presented in (b)-(d).

$$\text{res}_2 = \max\left\{ \frac{\left\|[\![A^p, B^p, C^p]\!] - [\![A^{p-1}, B^{p-1}, C^{p-1}]\!]\right\|_F}{\left\|[\![A^{p-1}, B^{p-1}, C^{p-1}]\!]\right\|_F}, \frac{\|\lambda^p - \lambda^{p-1}\|_F}{\|\lambda^{p-1}\|_F} \right\}. \tag{14}$$

The algorithm converges when $\text{res}_1 < \epsilon$ and $\text{res}_2 < \epsilon$, with typical $\epsilon = 10^{-3}$. The relative residual res_2 is inspired by the primal and dual residuals in [4]. Note that we are not seeking a *unique* decomposition of A, B, C, but rather a stable low rank construction $[\![A, B, C]\!]$ of observation $\mathcal{O}$ that has minimum sparse error.

3 Experiments on Synthetic Data

We test the NOTF algorithm on a synthetic dataset constructed as follows: (1) We create random sparse matrices $A \in \mathbb{R}^{50\times 3}, B \in \mathbb{R}^{20\times 3}, C \in \mathbb{R}^{10\times 3}$, where the sparse ratios (ratio of zero to non-zero values) are $70.67\%, 55\%$ and 30%, respectively. Each nonzero value is uniformly sampled from range $(0, 1)$. (2) Create the ground truth low rank matrix $\mathcal{X} = \mathcal{I}_{\{z:|z|>0\}}([\![A, B, C]\!]) \in \mathbb{R}^{50\times 20\times 10}$; it is the indicator tensor of nonzero values in the CP reconstruction from A, B, C; $[\![A, B, C]\!]$ has CP rank 3. Note that $\mathcal{X}$ is a discrete (binary) valued tensor with sparsity ratio 75.75%. (3) Create the observation tensor $\mathcal{O}$ by adding noise that flips a small portion of the binary values in $\mathcal{X}$.

NOTF reconstructs a low rank matrix $[\![\hat{A}, \hat{B}, \hat{C}]\!]$ with at most CP rank R, where $A \in \mathbb{R}^{50\times R}, B \in \mathbb{R}^{20\times R}, C \in \mathbb{R}^{10\times R}$. We vary the ratio of noise, i.e., the percentage of flips in $\mathcal{O}$, up to 10%, and the rank parameter R up to 10. We report on convergence iterations, false positive and false negative counts when reconstructing the binary values in tensor $\mathcal{X}$ and $\mathcal{O}$, and the mean square error for reconstructing $\mathcal{X}$ and $\mathcal{O}$. We compare NOTF with ℓ_0 norm in (1) with non-negative tensor factorization (NTF) baselines using the ℓ_1 norm and ℓ_2 norms. Both NOTF and baseline methods are initialized with a CP decomposition of the observation tensor $\mathcal{O}$, $\lambda^0 = 0$, with penalty parameter $\tau = 10$, and implementated in Matlab using the Tensor toolbox [3].

Fig. 1 presents the results of varying the noise ratio up to 10% (top) with rank parameter $R = 3$, and varying the rank parameter R up to 10 (bottom) with noise of 10%, for the synthetic dataset. We observe that the proposed NOTF solution with ℓ_0 norm performs well on the discrete measures (false positive and false negative counts in (b) and (c)); we consider non-zeros as positives and zeros as negatives in a tensor. NOTF achieves zero false positives and a relatively low false negative count over all noise ratios. The ℓ_2 baseline achieves zero false negatives, but the false positive counts are quite large. The ℓ_1 baseline achieves larger errors than NOTF on both positives and negatives, and is slower to converge.

Fig. 1 (d) shows that NOTF with ℓ_0 norm does not outperform the baselines for mean square error; this is not surprising. The discrete measurements are more important when reconstructing an occurrence

tensor. The error between the recovered low rank tensor $[\![\hat{A}, \hat{B}, \hat{C}]\!]$ and the observation $\mathcal{O}$ grows with the noise ratio, while the error between $[\![\hat{A}, \hat{B}, \hat{C}]\!]$ and the groundtruth $\mathcal{X}$ is relatively stable. This suggests that the recovered tensor $[\![\hat{A}, \hat{B}, \hat{C}]\!]$ can be used to de-noise the observation.

In Fig. 1 (bottom), we set a 10% noise ratio and vary the rank parameter R for the recovered tensor $[\![\hat{A}, \hat{B}, \hat{C}]\!]$. Note that R is an upper bound of the CP rank of $[\![\hat{A}, \hat{B}, \hat{C}]\!]$. NOTF could achieve zero for both false positives and false negatives when $R = 6$, which means that the ground truth can be completely recovered from the noisy observation. However, NOTF becomes unstable with larger R and leads to large false negative counts. A possible reason is that the initialization by CP decomposition of $\mathcal{O}$ becomes less stable when large R is used, and the least squares in (7)-(12) are often ill-posed and hard to solve. We finally observe that modeling the sparse error by the ℓ_0 norm brings an additional benefit in that the recovered $\hat{A}, \hat{B}$, and $\hat{C}$ are sparse; this leads to a clearer interpretation of each rank-one tensor as a community.

4 Experiments on the resMBS dataset

We explore the roles played by financial institutions (FIs) across multiple contracts (FCs) using NOTF with the ℓ_0 norm. ResMBS[5, 20, 19] contains extracted relationship of FI (e.g., Bank of America) playing a role (e.g., issuer) for a specific financial contract. The discrete values of occurrence tensor $\mathcal{O} \in \mathbb{R}^{971 \times 85 \times 27}$ indicate the counts of extractions of the specific (FC, FI, Role) occurrence from documents issued in 2005. $\mathcal{O}$ is sparse (1.02% non-zero values) and extraction noise is estimated to be $\leq 0.2\%$. We describe some observations here and present the relevant figures in Section 6 due to space limitations.

We vary the CP rank parameter R and reconstruct tensors for both the discrete observation $\mathcal{O}$ and its binary version. Performance is similar for both while it is notably slower to reconstruct discrete values (Fig. 2).We note that resMBS is challenging as the tensor is sparse. The false positives are relatively stable while the false negatives decrease as R increases. With $R = 20$, the total error count between the reconstructed tensor and the noisy observation is 3002; this roughly matches the expected errors of the information extractor. The histogram (Fig. 3 (left)) shows that errors for each FC is in a reasonable range (0, 20) with a mean of 3.

At last, we examine the discovered communities by NOTF for resMBS. Each rank-one tensor $a_r \circ b_r \circ c_r, r = 1, \ldots, R$ represents a community. Fig. 3 (right) presents the nonzero ratio and Fig. 4 presents the distribution of the reconstructed tensor component $\hat{A} = [a_r], \hat{B} = [b_r], \hat{C} = [c_r]$. An interesting observation is that the communities are "centered" around FIs, i.e., each community only contains one or two FIs. Some FIs could play various roles and appear in various FCs, while some FIs only play a limited number of roles in a limited number of FCs.

5 Discussion and future work

We present non-negative occurrence tensor factorization (NOTF) for analyzing heterogeneous networks. CP tensor decomposition is adapted to discover the embedded communities. The ℓ_0 norm is used to model the discrete tensor values and sparse errors, and the objective is solved with an efficient splitting optimization algorithm. NOTF is applied to both synthetic data and a new heterogeneous bipartite graph, resMBS, representing financial role relationships extracted from financial contracts. Preliminary results are promising and suggest that NOTF can be used to de-noise the occurrence tensors and identify communities in resMBS.

There are several directions for future work. The ℓ_0 norm is known to be difficult to optimize. The ℓ_p norm ($0 < p < 1$) satisfies the KL inequality, is often used as a surrogate, and may provide a theoretical convergence guarantee. The penalty parameter τ is crucial for both convergence speed and solution quality for nonconvex problems; adaptive ADMM[22, 21] which automates the selection of τ, achieves promising practical performance. To deal with the high sparsity of the resMBS tensor, domain-specific constraints (e.g., *each FC should contain an FI play role "Issuer"*) may boost performance. Finally, it is interesting to apply NOTF for analyzing some other heterogeneous networks that could be represented with an occurrence tensor.

Acknowledgments

ZX and TG were supported by US NSF grant CCF-1535902 and by US ONR grant N00014-15-1-2676. ZX and LR were supported by NSF grants CNS1305368 and DBI1147144, and NIST award 70NANB15H194.

References

[1] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15(1):2773–2832, 2014.

[2] A. Anandkumar, P. Jain, Y. Shi, and U. Niranjan. Tensor vs matrix methods: Robust tensor decomposition under block sparse perturbations. *arXiv preprint*, 2015.

[3] B. W. Bader and T. G. Kolda. Algorithm 862: MATLAB tensor classes for fast algorithm prototyping. *ACM Transactions on Mathematical Software*, 32(4):635–653, December 2006. doi: 10.1145/1186785.1186794.

[4] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. and Trends in Mach. Learning*, 3:1–122, 2011.

[5] D. Burdick, S. De, L. Raschid, M. Shao, Z. Xu, and E. Zotkina. resMBS: Constructing a financial supply chain graph from financial prospecti. In *SIGMOD DSMM workshop*. ACM, 2016.

[6] J. D. Carroll and J.-J. Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of "eckart-young" decomposition. *Psychometrika*, 35(3):283–319, 1970.

[7] X. Chen, Z. Han, Y. Wang, Q. Zhao, D. Meng, and Y. Tang. Robust tensor factorization with unknown noise. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5213–5221, 2016.

[8] D. Goldfarb and Z. Qin. Robust low-rank tensor recovery: Models and algorithms. *SIAM Journal on Matrix Analysis and Applications*, 35(1):225–253, 2014.

[9] T. Goldstein and S. Osher. The split Bregman method for L1-regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009.

[10] T. Goldstein, B. O'Donoghue, S. Setzer, and R. Baraniuk. Fast alternating direction optimization methods. *SIAM Journal on Imaging Sciences*, 7(3):1588–1623, 2014.

[11] Q. Gu, H. Gui, and J. Han. Robust tensor decomposition with gross corruption. In *Advances in Neural Information Processing Systems*, pages 1422–1430, 2014.

[12] R. A. Harshman. Foundations of the parafac procedure: Models and conditions for an" explanatory" multi-modal factor analysis. 1970.

[13] H. Huang and C. Ding. Robust tensor factorization using r 1 norm. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

[14] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

[15] A. P. Liavas and N. D. Sidiropoulos. Parallel algorithms for constrained tensor factorization via alternating direction method of multipliers. *IEEE Transactions on Signal Processing*, 63(20):5450–5463, 2015.

[16] K. Maruhashi, F. Guo, and C. Faloutsos. Multiaspectforensics: Pattern mining on large-scale heterogeneous networks with tensor analysis. In *Advances in Social Networks Analysis and Mining (ASONAM), 2011 International Conference on*, pages 203–210. IEEE, 2011.

[17] E. E. Papalexakis, C. Faloutsos, and N. D. Sidiropoulos. Tensors for data mining and data fusion: Models, applications, and scalable algorithms. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8 (2):16, 2016.

[18] A. Schein, J. Paisley, D. M. Blei, and H. Wallach. Bayesian poisson tensor factorization for inferring multilateral relations from sparse dyadic event counts. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1045–1054. ACM, 2015.

[19] Z. Xu and L. Raschid. Probabilistic financial community models with latent dirichlet allocation for financial supply chains. In *SIGMOD DSMM workshop*. ACM, 2016.

[20] Z. Xu, D. Burdick, and L. Raschid. Exploiting lists of names for named entity identification of financial institutions from unstructured documents. *arXiv preprint arXiv:1602.04427*, 2016.

[21] Z. Xu, S. De, M. A. T. Figueiredo, C. Studer, and T. Goldstein. An empirical study of admm for nonconvex problems. In *NIPS workshop on nonconvex optimization*, 2016.

[22] Z. Xu, M. A. Figueiredo, and T. Goldstein. Adaptive ADMM with spectral penalty parameter selection. *arXiv preprint arXiv:1605.07246*, 2016.

6 Appendix: experimental results for resMBS

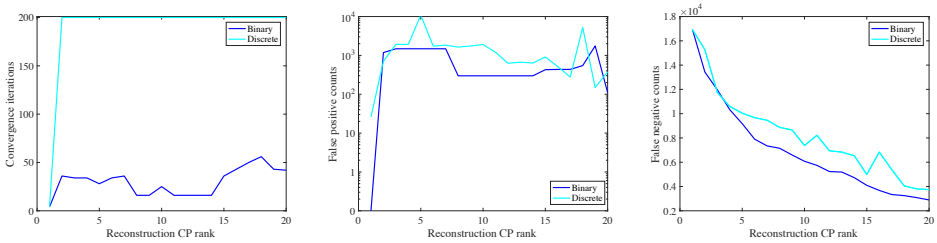

(a) (b) (c)

Figure 2: (a) Convergence iteration, (b) false positive counts, and (c) false negative counts when vary the tensor CP rank R for the reconstruction of resMBS dataset. Both binary and discrete tensor of resMBS are tested. Note that reconstruction errors presented in (b)(c) are based on noisy observation $\mathcal{O}$ as the groundtruth is unknown.

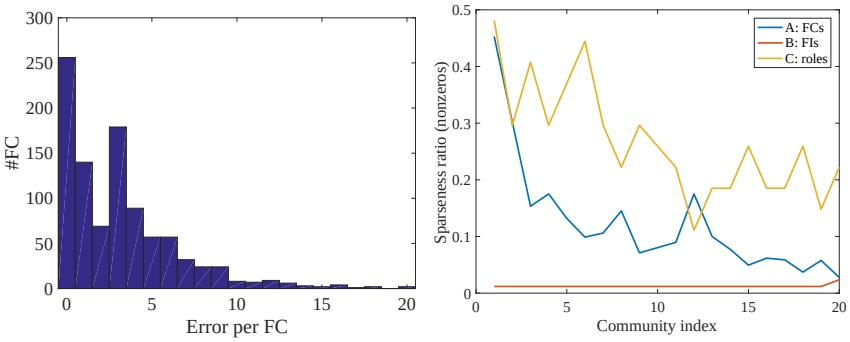

Figure 3: (left) Histogram of errors when constructed resMBS tensor with $R = 20$. (right) Nonzero ratio for the constructed tensor component $\hat{A}, \hat{B}, \hat{C}$. Each rank-one tensor $a_r \circ b_r \circ c_r, r = 1, \ldots, R$ represents a community.

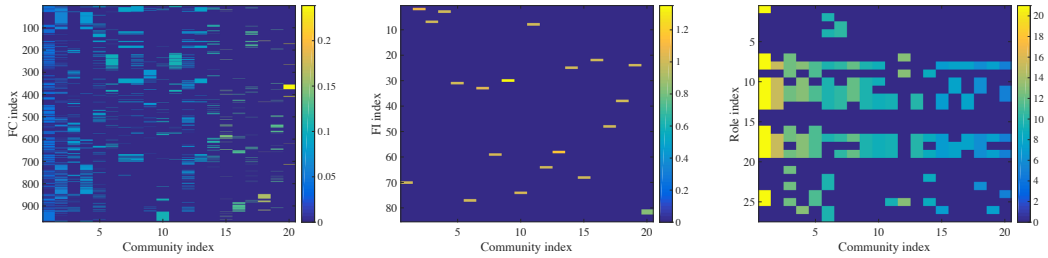

(a) (b) (c)

Figure 4: The distribution of FC, FI, and roles in each community.